IN THE CLAIMS:

This listing of claims will replace all prior versions, and listings, of claims in the application:

1. (Currently Amended) A method of scheduling the passage of data cells from [[M]] N low-bandwidth data sources to M low-bandwidth data destinations in a data switching apparatus which includes:

M/N ingress multiplexers, each arranged to receive data cells from a respective set of N said low-bandwidth data sources,

M/N egress multiplexers, each arranged to transmit data cells to a respective set of N said low-bandwidth data destinations,

a master control unit, and

a central switch having M/N high-bandwidth input ports arranged to receive data cells from respective said ingress multiplexers, and M/N high-bandwidth output ports arranged to transmit data cells to respective said egress multiplexers, the central switch selectively interconnecting the input ports and output ports, under the direction of the master control unit,

the method comprising the steps of:

maintaining N input queues in each of said ingress multiplexers for queuing data cells received from the N respective said data sources, and maintaining M virtual output queues for queuing data cells directed to respective said data destinations;

maintaining a respective ingress port table in each of said ingress multiplexers, each ingress port table having NxM entries, each entry corresponding to a respective combination of a said data source for that ingress port and a said data destination,

transferring data cells from said input queues to said virtual output queues by each of said ingress multiplexers with a relative frequency according to a value of the corresponding entry of the ingress port table;

maintaining a respective egress port table in each of said ingress multiplexers, the egress port table, having M entries, each entry corresponding to a respective said data destination,

transferring data cells from said virtual output queues to said respective input ports of the central switch by each of said ingress multiplexers with a relative frequency according to the value of the corresponding entry of the egress port table;

maintaining a central allocation table in the master control unit having $(M/N)^2$ entries, each corresponding to a respective combination of an input port and an output port, and

controlling the central switch by the master control unit to interconnect pairs of said input ports and output ports with a relative frequency according to the value of the corresponding entry of the central allocation table;

whereby said ingress port tables, egress port tables and central allocation table together determine the bandwidth through the digital data switching apparatus from each said data source to each said data destination.

2. (Previously Presented) The method according to claim 1 wherein each said ingress multiplexer, for each virtual output queue, transfers data cells to that virtual output queue from said input queues in accordance with a N-way weighted round robin, using N weights determined respectively by the N entries of the ingress port table for that virtual output queue.

3. (Previously Presented) The method according to claim 2, wherein each weight is defined by a number of bits w, and the N-way weighted round robin for each virtual output queue is implemented by an $N(2^w-1)$-way unweighted round robin using a request vector list constructed by interleaving N words of (2^w-1) bits each, each word corresponding to a respective input queue and having a number of "1"s determined by the entry of the ingress port table for that input queue and that virtual output queue.

4. (Previously Presented) The method according to claim 3, wherein the request vector list is separated into a plurality of round robin blocks, each corresponding to a respective input queue, a first round robin process being performed independently within each block, and a second round robin process being performed to make a selection among the blocks.

5. (Previously Presented) The method according to claim 1, wherein the ingress port table, the egress port table and the central allocation table are all programmed from an external source.

6. (Previously Presented) The method according to claim 5, wherein the external source uses parameters characterizing the length of each virtual output queue and the urgency of each virtual output queue.

7. (Previously Presented) The method according to claim 6, wherein the external source uses a set of sensitivities relating to length, urgency and pseudo-static bandwidth allocation.

8. (Currently Amended) A digital data switching apparatus for transmitting data from [[M]] N low-bandwidth data sources to M low-bandwidth data destinations, comprising:

M/N ingress multiplexers for receiving data cells from respective sets of N said low-bandwidth data sources,

M/N egress multiplexers for transmitting data cells to respective sets of N said low-bandwidth data destinations,

a master control unit,

a central switch having M/N high-bandwidth input ports arranged to receive data cells from respective said ingress multiplexers, and M/N high-bandwidth output ports arranged to transmit data cells to respective said egress multiplexers, the central switch being arranged selectively to interconnect the input ports and output ports, under the direction of the master control unit,

each said ingress multiplexer being arranged to maintain N input queues for queuing data cells received from respective said data sources, and to maintain M virtual output queues for queuing data cells directed to respective said data destination;

wherein each ingress multiplexer is arranged to maintain a respective ingress port table, each ingress port table having NxM entries, each entry corresponding to a respective combination of a said data source and a said data destination, and each ingress multiplexer is arranged to transfer data cells from said input queues to said virtual output queues with a relative frequency according to value of the corresponding entry of the ingress port table;

each ingress multiplexer is arranged to maintain a respective egress ort table, the egress port table having M entries, each corresponding to a respective said data destination, and each ingress ~~multiplier~~ multiplexer is arranged to transfer data cells from said virtual output queues to said respective input ports of the central switch with a relative frequency according to value of the corresponding entry of the egress port table,

and the master control unit is arranged to maintain a central allocation table having $(M/N)^2$ entries, each corresponding to a respective combination of an input port and an output port, and the master control unit controls the central switch to interconnect pairs of said input ports and output ports with a relative frequency according to the value of the corresponding entry of the central allocation table;

whereby said ingress port tables, egress port tables and central allocation table together determine the bandwidth through the digital data switching apparatus from each said data source to each said data destination.

9. (Previously Presented) The apparatus according to claim 8, wherein each said ingress multiplexer is arranged, for each virtual output queue, to transfer data cells to that virtual output queue from said input queues in accordance with a N-way weighted round robin, using N weights determined respectively by the N entries of the ingress port table for that virtual output queue.

10. (Previously Presented) The apparatus according to claim 8, wherein each weight has a number of bits w, and the N-way weighted round robin for each virtual output queue is implemented by an $N(2^w-1)$-way unweighted round robin using a request vector list constructed by interleaving N words of (2^w-1) bits each, each word corresponding to a respective input queue and having a number of "1"s determined by the entry of the ingress port table for that input queue and that virtual output queue.

11. (Previously Presented) The apparatus according to claim 10, wherein the request vector list is separated into a plurality of round robin blocks, each corresponding to a respective input queue, each ingress multiplexer being arranged to perform a first round robin process independently within each block, and a second round robin process to make a selection among the blocks.

12. (Previously Presented) The apparatus according to claim 8, further comprising an external source unit arranged to program the ingress port table, the egress port table and the central allocation table.

13. (Previously Presented) The apparatus according to claim 12, wherein the external source unit is arranged to operate using parameters characterizing the length of the virtual output queue and the urgency of the virtual output queue.

14. (Previously Presented) The apparatus according to claim 13, wherein the external source unit is arranged to operate using a set of sensitivities relating to the length, urgency and pseudo-static bandwidth allocation.